

(12) UK Patent Application (19) GB (11) 2 358 261 (13) A

(43) Date of A Publication 18.07.2001

(21) Application No 0001030.8

(22) Date of Filing 17.01.2000

(71) Applicant(s)

ARM Limited
(Incorporated in the United Kingdom)
110 Fulbourn Road, Cherry Hinton, CAMBRIDGE,
CB1 9NJ, United Kingdom

(72) Inventor(s)

Edward Colles Nevill

(74) Agent and/or Address for Service

D Young & Co
21 New Fetter Lane, LONDON, EC4A 1DA,
United Kingdom

(51) INT CL⁷

G06F 9/42 9/455

(52) UK CL (Edition S)

G4A APB

(56) Documents Cited

None

(58) Field of Search

UK CL (Edition R) G4A APB
INT CL⁷ G06F
ONLINE:WPLEPODOC,JAPIO

(54) Abstract Title

Data processing with native and interpreted program instruction words

(57) An Interpreter invocation mechanism for switching between execution of native instruction words (ARM, Thumb) and Interpreted instruction words (Photon) uses a subroutine call instruction (14) to start execution of the interpreter. The return address of the subroutine call instruction (14) is used as an address pointer to the start of the interpreted code (16). The interpreted code may terminate with an Exit instruction whereupon normal native code execution resumes using the instruction at the immediately following memory address or alternatively with a Return instruction that recovers a return address previously stored to stack memory.

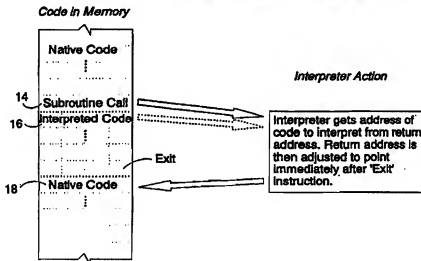


Fig.3

GB 2 358 261 A




Fig.1

1e






Fig. 2








Fig. 4

DATA PROCESSING WITH NATIVE AND INTERPRETED
PROGRAM INSTRUCTION WORDS

5 This invention relates to the field of data processing systems. More particularly, this invention relates to data processing system operating under control of a mix of native and interpreted program instruction words.

Computer programs or subroutines written in the form of interpreted program instructions words are typically smaller in size than the same computer programs or subroutines written in the form of native programs instruction words. High code
10 density is desirable as it reduces system memory requirements. However, computer programs or subroutines written in the form of interpreted program instruction words are typically slower to execute than the same computer programs or subroutines written in the form of native program instruction words due to the additional processing overhead associated with interpretation. The different characteristics of
15 computer programs or subroutines written in these two ways has the result that it is desirable to be able mix the two types of program words and rapidly and efficiently switch between them.

When executing a portion of a computer program in the form of native program instruction words and it is desired to switch to executing interpreted program
20 instruction words, then the instruction interpreter must be invoked and passed the location of the interpreted program instruction words. It is often not possible or desirable to add a special purpose native program instruction word to carry out this task and so several native program instruction words have to be used to perform the switch. The use of several native program words to perform the switch is
25 disadvantageously slow and uses extra memory capacity.

Viewed from one aspect the invention provides apparatus for processing data under control of a computer program, said apparatus comprising:

- a processing unit responsive to native program instruction words to perform data processing operations;
- 30 an instruction interpreter responsive to one or more interpreted instruction words specifying a data processing operation to execute native instruction words upon said processing unit to perform said data processing operation; and

a memory for storing said computer program; wherein

said computer program includes both native instruction words and interpreted instruction words;

a native code portion invokes interpretation of an interpreted code portion by
5 executing a native code call instruction to said instruction interpreter;

execution of said native code call instruction triggers generation of a return address specifying a location within said memory for said native code call instruction; and

said instruction interpreter uses said return address as a pointer to said interpreted
10 code portion within said memory.

Using a native code call instruction, as is normally provided in the instruction set, to call/invoke the interpreter avoids the need for a special purpose native program instruction word to perform the switch. The location of the interpreted program instruction words is passed to the interpreter in the return address of the native code
15 call instruction and so only a single native code instruction is needed to make the switch.

Whilst it is possible to do otherwise, it is preferred that said return address is an address immediately succeeding an address of said native code call instruction within said memory.

20 As the return address points to the interpreted program instruction words, then these will normally be provided immediately following the native code call instruction.

In order to manage the switch back to native code program instruction words, one preferred arrangement is that said instruction interpreter writes a new value of
25 said return address for use by a native code return instruction as a pointer to a next native code instruction to be executed when interpretation of said interpreted code portion is finished and return is being made to execution of said next native code portion of said computer program.

In accordance with the above, the return to native code execution may be
30 made by an arrangement that said interpreted code portion finishes with an interpreted

return instruction that returns execution to a native code portion at an address pointed to by said return address.

As an alternative, in some situations it may be preferred that said interpreted code portion finishes with an interpreted exit instruction that results in execution of a next native code portion stored within said memory at an address immediately following said exit instruction without a return being made corresponding to said native code call instruction.

It is common for call return addresses to be stored in stack memory, however in preferred embodiments said processing unit includes a bank of data processing registers and said return address is stored within a predetermined data processing register within said bank of data processing registers upon execution of said native code call instruction.

Having the return address in a register allows use and manipulation of this address to take place without the overhead of having to recover the address from stack memory.

Whilst having the return address available in a register for immediate use is an advantage, it is also sometimes desirable to release that register for other uses or to provide longer term storage of the return address. Accordingly, in preferred embodiments of the invention said memory includes a stack memory region and said return address is copied from said predetermined data processing register to said stack memory region upon invocation of said interpreted code portion.

The interpreter can be made more simple, compact and fast in preferred embodiments in which each interpreted code instruction word corresponds to a native code instruction word.

In practice the efficiency of the interpreter can be improved and the code density increased in preferred embodiments in which said native code instruction words form a native code instruction set and said interpreted code instruction words form an interpreted code instruction set, data processing operations provided by said interpreted code instruction set being a subset of data processing operations provided by said native code instruction set.

The interpretation simplicity and operational efficiency are further improved in embodiments in which said native code instruction words are X-bit instruction

in response to one or more interpreted instruction words specifying a data processing operation, executing with an instruction interpreter native instruction words upon said processing unit to perform said data processing operation; and

storing said computer program in a memory; wherein

5 said computer program includes at least one native code portion and at least one interpreted code portion;

a native code portion invokes interpretation of an interpreted code portion by executing a native code call instruction to said instruction interpreter;

10 execution of said native code call instruction triggers generation of a return address specifying a location within said memory for said native code call instruction; and

said instruction interpreter uses said return address as a pointer to said interpreted code portion within said memory.

15 An embodiment of the invention will now be described, by way of example only, with reference to the accompanying drawings in which:

Figure 1 is a schematic diagram of a data processing system;

Figure 2 schematically illustrates switching from native code to interpreted code;

20 Figure 3 illustrates a first mechanism for returning to native code execution from interpreted code execution; and

Figure 4 illustrates an alternative way of invoking execution of interpreted code and a second mechanism for returning to native code execution from interpreted code execution.

25 Figure 1 illustrates a data processing system 2 comprising an integrated circuit 4 coupled to a main memory 6. Within the integrated circuit 4 there is a processor core 8 and a cache memory 10. Within the processor core 8 there is an instruction decoding and control unit 11 coupled to a register bank 12. It will be appreciated that Figure 1 is a highly schematic representation of a data processing system that will in practice contain many more parts.

30 The main memory 6 stores software programs and data to be processed. The processor core 8 fetches instruction words and data words from the main memory.

BN
0000
10 00

102348
M

15

14

13

12

6

11

9

8

10

5

4

3

2

1

number of photon instruction words can be reserved for controlling the interpreter and other uses which do not have direct single equivalents in the native instruction words.

5 Providing substantially each of the Photon instruction words in a one-to-one correspondence with native instruction words simplifies the required interpreter such that this interpreter consumes relatively few resources and is able to operate rapidly and predictably. The interpreter may be a software interpreter or a hardware interpreter. Using the Photon instruction set, a hardware interpreter will have the advantage of requiring comparatively less storage space and a software interpreter will have the advantage of requiring comparatively fewer circuit elements.

10 The software interpreter is "locked-down" within the cache memory 10. Thus, when the software interpreter is first invoked it will be loaded into the cache memory 10 from the main memory 6 and thereafter remain within the cache memory 10 to be available for high speed use. Alternatively, the software interpreter may be pre-loaded into the cache memory 10 from the main memory 6 when the data processing
15 system 2 is initialized. In either case, having a relatively small and efficient software interpreter that may be wholly stored within the cache memory 10 without consuming too great a proportion of the cache memory 10 is a strong performance advantage.

The software interpreter itself is formed of native instruction words and serves to parse the interpreted instruction words and map them to corresponding native
20 instruction words. It will be appreciated that whilst substantially each Photon instruction word corresponds to an ARM instruction, the software interpreter may need to execute a few ARM instructions to set up the data in the correct registers and form prior to executing the ARM instruction that corresponds to the interpreted Photon instruction.

25 In order that the system may operate effectively, it is important to have an efficient mechanism for switching between the execution of native instruction words and the execution of interpreted instruction words. The native instruction words instruction set is often already fixed and so it is not possible to add a special purpose switching instruction to the native instruction set. Furthermore, bit-space within a
30 native instruction set is often a critical resource and so adding a switching instruction would reduce the amount of this resource available for other uses.

Figure 2 illustrates a technique for switching between execution of native instruction words and interpreted instruction words. The native instruction words include a subroutine call instruction that is normally used to jump to native code subroutines. When such a subroutine call is made, hardware mechanisms are already
 5 in place that store a return address for the call such that native code execution can be resumed at the instruction following the subroutine call once the subroutine has completed. The interpreter invocation mechanism illustrated in Figure 2 uses a subroutine call 14 from the native code to the software interpreter code. The interpreted code 16 is stored in the memory locations immediately following the
 10 subroutine call. Accordingly, the return address stored in response to the subroutine call 14 corresponds to the memory location of the start of the interpreted code 16. Accordingly, the software interpreter reads the return address and uses this as a pointer to the interpreted code 16 which it is desired to execute. It will be seen from Figure 2 that the native code comprises 32-bit native instruction words and the
 15 interpreted code comprises 8-bit interpreted instruction words, such that four interpreted instruction words can be stored in the same space that is required for a single native code instruction word. A single subroutine call instruction 14 is needed to perform the required switch to the start of execution of interpreted instruction words. The subroutine call instruction 14 is already provided within the native code
 20 instruction set and so no additional bit-space within the native code instruction set is consumed. Furthermore, if the native code instruction set is one that is already established, then no modifications need to be made to the native code instruction set or hardware to utilize the interpreted code invocation technique illustrated in Figure 2.

Figure 3 illustrates a first mechanism for returning to native code execution
 25 from interpreted code execution. A special purpose Exit instruction is provided within the interpreted code instruction set. The software interpreter is responsive to such an Exit instruction to simply cease interpretation of instruction words and instead pass control to the immediately following next native code instruction 18. Using this technique interpreted code can be freely mixed in with native code.

30 Figure 4 illustrates an alternative technique for invoking the software interpreter and a second mechanism for returning to native code execution from interpreted code execution. Whilst Figure 3 allows interpreted code to be embedded within native code, it is not well suited to situations in which the interpreted code is

genuinely a subroutine that it is desired to invoke from several different points within the native code of the program as a whole. Figure 4 illustrates a solution to this problem. A subroutine call 14 is issued to a subroutine Routine1 in memory that appears to the overall native code program to be a native code subroutine in the normal way. The subroutine Routine1 stores the return address to the native code instruction 18 by pushing it to a stack memory area with native code instruction 20. The subroutine Routine1 then makes a second subroutine call 22 to the start address of the interpreter software. This second subroutine call invokes interpreted code execution in the same way as illustrated in Figure 2 in that the return address from the subroutine call instruction 22 is used as a pointer to the start of the interpreted code 16.

At the end of the interpreted code there is a return instruction that is acted upon by the software interpreter to recover the return address from the subroutine call instruction 14 that was stored in the stack memory area. The software interpreter then passes control using this recovered return address to the native code instruction 18 that is pointed to by the return address.

In this way, the interpreted code can effectively be treated as a normal native code subroutine by the rest of the native code program.

An alternative way of viewing the invention is set out as follows.

This describes a mechanism for invoking an interpreter to execute interpreted code. The mechanism uses a subroutine call followed immediately by the code to be interpreted. The interpreter uses the return address of the subroutine as the starting address of the code to be interpreted. If the interpreter needs to return from interpreting code it adjusts the return address to point to the end of the interpreted code so that on return from the interpreter, normal (native) program execution will continue at the instruction after the interpreted code.

Alternatively, if the interpreter needs to perform a return directly to the caller of the original function containing the interpreted code, it can pop the return frame from the stack. Execution will then continue at the instruction after the call to the function containing the interpreted code.

A code example of the mechanism illustrated in Figure 2 using the ARM processor with the Photon interpreter is as follows:

```

CODE32                                ; Instruct the assembler that what follows is
                                      ARM code.

BL      Execute_Photon                ; Call the Photon interpreter
5  PHOTON                             ; Instruct the assembler that what follows is a
                                      Photon byte code

<Photon byte code inserted here>

```

10 In the above code example the processor executes the first instruction (BL) as a native ARM instruction. The BL instruction (a subroutine call in ARM) invokes the interpreter which then executes then native ARM instructions necessary to interpret the bytes following the BL instruction as Photon byte code.

Benefits

15 Previously to invoke the interpret on a section of interpreted code the processor would have to load the address of the code to be executed into a register (say R0) and then call the interpreter. e.g.

```

ADR      r0, Interpreted_code
BL      Execute_Photon
20  ...

```

Interpreted_Code

<Interpreted code here>

25 With the "Interpreter Invocation Mechanism" the code to be interpreted can be placed immediately after the call to Execute_Photon. There is then no need to load R0 with the address of the code to be interpreted. This has the following benefits.

- There is a saving of one instruction per interpreter invocation

- There is no need to reserve a register to hold the address of the code to be interpreted. As the point of a subroutine call is often the point where the register pressure is greatest there may not normally be a free register available to use to point to the interpreted code. For example, in the above
5 R0 is typically used to hold the 1st argument to the function.

A code example of the mechanism illustrated in Figure 3 using the ARM processor with the Photon interpreter is as follows:

```

CODE32                                ;Instruct the assembler that what follows
10                                is ARM code
BL      Execute_Photon                ;Call the Photon interpreter
PHOTON                                ;Instruct the assembler that what follows is a
                                      Photon byte code
<Photon byte code inserted here>
15 EXIT                              ;Photon Exit interpreter instruction
CODE32                                ;Instruct the assembler that what follows is
                                      ARM code
25

```

As in the first example the processor executes the first 2 instructions as native
20 ARM code and then invokes the interpreter to execute Photon code immediately following the BL.

Execution of Photon code by the interpreter continues until the interpreter encounters a Photon 'EXIT' instruction. The interpreter then takes its IP register (Interpreter Pointer) and uses this to return to the ARM instructions immediately
25 following the Photon code.

Benefits

The use of an EXIT instruction within the interpreter allows the interpreter to return directly to the code after the interpreted code. This allows interpreted code to be intermingled with native code. The possible benefits of this are.

- A compiler can choose to compile sections of C (or other high level language code) to either native machine code, or interpreted code depending on which is more efficient in terms of code size or performance

In a further example a subroutine call is made to a 1st routine which is part of its entry sequence performs a subroutine call to invoke the interpreter. The interpreter then executes interpreted instructions in the 1st routine until it encounters an interpreted RETURN instruction. Rather than continuing execution after the RETURN instruction as in the example of the EXIT instruction above it instead returns immediately to the caller of the 1st routine.

- 10 A code example of the mechanism illustrated in Figure 4 using the ARM processor with the Photon interpreter is as follows:

```

CODE32                                ;Instruct the assembler that what follows
                                      is ARM code

15  <ARM code>

    BL      Routine1                  ;Ordinary call to Routine1

    <ARM code>                        ;Execution continues here after execution of
                                      Routine1

;-----

20  CODE32                                ;Instruct the assembler that what follows is
                                      ARM code

    STMDB   sp!, {sp, lr}             ;Save return address for this function and also
                                      save stack pointer to allow for stack unwinding.

    BL      Execute_Photon            ;Call the Photon interpreter

25  PHOTON                                ;Instruct the assembler that what follows is a
                                      Photon byte code

    <Photon byte code inserted here>

    RETURN                            ;Photon RETURN interpreter instruction
  
```

In this example an initial subroutine call is made to Routine1. Routine1 then performs an entry sequence which consists of a STMDB instruction which saves the link register (LR) and the stack pointer (SP). The link register contains the return address for the Routine.

- 5 Execution of a Photon code by the interpreter continues until the interpreter encounters a Photon RETURN instruction. The interpreter then pops the link register and stack pointer from the stack and moves the link register into the program counter (PC) by performing the following sequence

```
10  LDMIA    sp, {sp, lr}
    BX      lr
```

This has the effect of returning to the instruction immediately following the original subroutine call to Routine1.

15 Benefits

The use of the Interpreter Invocation Mechanism in the entry sequence of a function combined with the RETURN interpreted instruction allows a function to be written in interpreted code but be transparently viewed from the rest of the system as though it were written in native code. This has the following benefits.

- 20 • A compiler can choose to compile entire functions in either native code or interpreted code depending on what is most efficient in terms of code density or performance without needing to know whether the target function is interpreted code or native code. Since, for languages such as C different functions may be compiled from separate source files at different
- 25 times the compiler can not know whether the target function is interpreted code or native code
- As software is usually developed in a modular fashion with different developers working on different sections of code it eases the task of software development as individual software developers can choose to write code in either interpreted code or native code and know that software
- 30 written by other developers can call their software regardless.

Further Benefits

The above examples show the benefits as applied to ARM and Photon. These benefits also apply to other architectures and other interpreted languages.

For example the benefits of the above would also apply to the Intel
5 80x86/Pentium architectures with Java as the interpreted language.

This would allow Java and C code to be intermixed freely. For example following a C Entry Point an Entry sequence could be executed followed by a call to a Java interpreter. A Java return instruction would return control to the instruction following the Java code.

10 Alternatively one could use the same idea with one of the 80x86/Pentium architectures with a BASIC languages such as Visual BASIC as the interpreted language.

CLAIMS

1. Apparatus for processing data under control of a computer program, said apparatus comprising:
- 5 a processing unit responsive to native program instruction words to perform data processing operations;
- an instruction interpreter responsive to one or more interpreted instruction words specifying a data processing operation to execute native instruction words upon said processing unit to perform said data processing operation; and
- 10 a memory for storing said computer program; wherein
- said computer program includes both native instruction words and interpreted instruction words;
- a native code portion invokes interpretation of an interpreted code portion by executing a native code call instruction to said instruction interpreter;
- 15 execution of said native code call instruction triggers generation of a return address specifying a location within said memory for said native code call instruction; and
- said instruction interpreter uses said return address as a pointer to said interpreted code portion within said memory.
- 20
2. Apparatus as claimed in claim 1, wherein said return address is an address immediately succeeding an address of said native code call instruction within said memory.
- 25 3. Apparatus as claimed in any one of claims 1 and 2, wherein said instruction interpreter writes a new value of said return address for use by a native code return instruction as a pointer to a next native code instruction to be executed when interpretation of said interpreted code portion is finished and return is being made to execution of said next native code portion of said computer program.

30

4. Apparatus as claimed in claim 3, wherein said interpreted code portion finishes with an interpreted return instruction that returns execution to a native code portion at an address pointed to by said return address.

5. Apparatus as claimed in any one of claims 1 and 2, wherein said interpreted code portion finishes with an interpreted exit instruction that results in execution of a next native code portion stored within said memory at an address immediately following said exit instruction without a return being made corresponding to said native code call instruction.

6. Apparatus as claimed in any one of the preceding claims, wherein said processing unit includes a bank of data processing registers and said return address is stored within a predetermined data processing register within said bank of data processing registers upon execution of said native code call instruction.

7. Apparatus as claimed in claim 6, wherein said memory includes a stack memory region and said return address is copied from said predetermined data processing register to said stack memory region upon invocation of said interpreted code portion.

8. Apparatus as claimed in any one of the preceding claims, wherein substantially each interpreted code instruction word corresponds to a native code instruction word.

9. Apparatus as claimed in claim 8, wherein said native code instruction words form a native code instruction set and said interpreted code instruction words form an interpreted code instruction set, data processing operations provided by said interpreted code instruction set being a subset of data processing operations provided by said native code instruction set.

10. Apparatus as claimed in any one of the preceding claims, wherein said native code instruction words are X-bit instruction words and said interpreted code instruction words are Y-bit instruction words, X being greater than Y.
- 5 11. Apparatus as claimed in any one of the preceding claims, wherein said native code instruction words specify N-bit data processing operations to be performed by said processing unit and said interpreted code instruction words specify N-bit data processing operations to be performed by said processing unit.
- 10 12. Apparatus as claimed in claims 10 and 11, wherein X is equal to N.
13. Apparatus as claimed in claim 12, wherein Y is $N/4$.
14. Apparatus as claimed in any one of the preceding claims, wherein said
15 instruction interpreter is a software interpreter provided by an interpreter computer program formed of native code instruction words executed by said processing unit.
15. Apparatus as claimed in claim 14, comprising a cache memory and wherein said software interpreter is, in use, stored within said cache memory.
- 20 16. Apparatus as claimed in claim 15, wherein said software interpreter is stored within said cache memory both when executing native program instruction words and interpreted program instruction words.
- 25 17. Apparatus as claimed in any one of the preceding claims, wherein portions of said computer program having less than a threshold required execution speed are provided as interpreted code instructions words and portions of said computer program having greater than said threshold required execution speed are provided as native code instructions words.
- 30

18. Apparatus as claimed in any one of the preceding claims, wherein portions of said computer program having less than a threshold execution frequency are provided as interpreted code instructions words and portions of said computer program having greater than said threshold execution frequency are provided as native code instructions words.

19. A method of processing data under control of a computer program, said method comprising the steps of:

in response to native program instruction words, performing data processing operations with a processing unit;

in response to one or more interpreted instruction words specifying a data processing operation, executing with an instruction interpreter native instruction words upon said processing unit to perform said data processing operation; and

storing said computer program in a memory; wherein

- said computer program includes at least one native code portion and at least one interpreted code portion;

a native code portion invokes interpretation of an interpreted code portion by executing a native code call instruction to said instruction interpreter;

- execution of said native code call instruction triggers generation of a return address specifying a location within said memory for said native code call instruction; and

said instruction interpreter uses said return address as a pointer to said interpreted code portion within said memory.

20. Apparatus for processing data substantially as hereinbefore described with reference to the accompanying drawings.

21. A method of processing data substantially as hereinbefore described with reference to the accompanying drawings.



Application No: GB 0001030.6
Claims searched: 1-21

Examiner: Mike Davis
Date of search: 24 August 2000

Patents Act 1977
Search Report under Section 17

Databases searched:

UK Patent Office collections, including GB, EP, WO & US patent specifications, in:

UK Cl (Ed.R): G4A (APB)

Int Cl (Ed.7): G06F

Other: Online: WPI, EPODOC, JAPIO

Documents considered to be relevant:

Category	Identity of document and relevant passage	Relevant to claims
	None	

X Document indicating lack of novelty or inventive step	A Document indicating technological background and/or state of the art
Y Document indicating lack of inventive step if combined with one or more other documents of same category.	P Document published on or after the declared priority date but before the filing date of this invention.
& Member of the same patent family	E Patent document published on or after, but with priority date earlier than, the filing date of this application.

An Executive Agency of the Department of Trade and Industry